

Overview of Diagnostic Methods

Philip Haves
Lawrence Berkeley National Laboratory

Summary

This paper reviews the use of diagnostics in buildings, both in commissioning and in operations and maintenance. The difference between *fault detection* and *fault diagnosis* is defined. The kinds of fault that can arise at different stages of the building life cycle are described and the ways that these different kinds of fault can be detected are discussed. Other issues addressed include manual and automatic methods and whole building vs. component-level approaches.

Introduction

Many buildings fail to perform as well as expected because of problems that arise at various stages of the life cycle, from design planning to operation. These problems relate to different parts of the building, including the envelope, the HVAC system and the lighting system. The consequences include increased energy costs, occupant discomfort, poor productivity and health problems and increased maintenance costs. Manual methods for commissioning and trouble-shooting are available and are starting to be used more widely. Over the last few years, a number of researchers in the US and elsewhere have developed automated methods for diagnosing faults in buildings, primarily in HVAC systems (Hyvarinen 1997) and are now starting to demonstrate their use in real buildings. What follows is an overview of the application of diagnostic methods in buildings. One aim of this overview is to provide a conceptual framework and a common set of terms for describing diagnostic methods.

Definition of ‘Diagnostics’

There are two levels or stages of ‘diagnostics’:

- fault detection
- fault diagnosis

Fault *detection* is the determination that the operation of the building is incorrect or unacceptable in some respect. This can be done either by:

- assessing the performance of all or part of the building over a period of time (e.g. from utility bills or complaint frequencies) and then comparing that performance to what is expected;

- monitoring the temperatures, flow rates and power consumptions and continuously checking for incorrect operation or unsatisfactory performance.

Unacceptable behavior may occur over the whole operating range or be confined to a limited region and hence only occur at certain times.

Fault *diagnosis* is the identification or localization of the cause of faulty operation. It involves determining which of the possible causes of faulty behavior are consistent with the observed behavior. Automated fault diagnosis relies entirely on sensors and so may not be able to identify the nature of the fault unambiguously, although it may be able to eliminate some of the possible causes.

In addition to establishing the physical cause of faulty operation (e.g. fouled coil, overridden control function), it is desirable to be able to estimate both the cost of fixing the fault and the cost of delaying fixing it.

Types of Fault

Faults can arise at every stage in the building life cycle:

- planning, e.g. actual building occupancy incorrectly anticipated and specified;
- design – design basis fails to meet design specification, e.g. undersizing;
- incorrect equipment supplied, e.g. inappropriate substitution;
- incorrect installation, e.g. motor phases wired incorrectly;
- malicious damage – most likely during construction but possible at any stage;
- testing and balancing (TAB), e.g. incorrect balancing of air flows to different zones;
- commissioning, e.g. poor control loop tuning;
- changes made by the operator that may improve operation under a limited set of conditions, e.g. disabling mixing box dampers;
- equipment failure, e.g. broken fan belt;
- equipment degradation, e.g. coil fouling.

In order to be able to detect a fault, i.e. to know that the observed operation is incorrect or unsatisfactory, it is necessary to have some reference, or baseline, to which actual performance may be compared. Possible sources of this baseline include:

- design brief, e.g. provide comfort at given occupancy level
- codes and standards, e.g. California's Title 24;
- rules of thumb – 'standard practice';
- design basis, e.g. heating coil should provide a particular duty under particular conditions;
- manufacturer's specification, e.g. chiller performance map;
- performance of comparable buildings, e.g. whole building energy consumption (BTU/sq.ft/yr)

- in situ test performance, e.g. results of acceptance test during commissioning
- ‘normal’ operating data, collected when system is assumed to be operating correctly

The choice of baseline depends on the kinds of faults that need to be detected. For example, comparing the current capacity of a cooling coil to the capacity observed during commissioning can show that it is fouled but cannot show that the wrong coil was selected. The earlier in the life-cycle a particular fault arises, the more restricted is the choice of baseline that can be used to detect the problem, since the baseline has to originate earlier in the life-cycle than the fault in order for it to be detectable.

Some baselines are easier to establish than others. For example, rules-of-thumb can be used to provide an approximate baseline for certain aspects of performance. In general, more accurate baselines are more expensive to obtain since they are specific to the particular installation.

Physical faults in HVAC systems can be divided into two classes: *abrupt* faults, e.g. broken fan belt, and *degradation* faults, e.g. coil fouling. Abrupt faults are easier to detect, since they generally result in a sudden failure of some part of the plant, although they are not necessarily easier to diagnose. In the case of degradation faults, it is necessary to define a threshold, below which the fault is considered insignificant and above which it is considered desirable to detect the fault. This threshold may not be easy to set. In principle, this threshold should be determined by some kind of cost-benefit analysis. A low threshold can allow the detection of smaller faults and the earlier detection of major faults but requires better sensors and better detection methods to avoid false alarms.

Performance Verification and Performance Monitoring

Diagnostic methods can be applied either during on-going operation or at specific critical phases in the building life cycle, e.g. commissioning, retrofit or change of ownership. Short-term testing, sometimes known as *performance verification*, involves deliberate exercising of the appropriate systems, e.g. by changing set-points or introducing artificial heat loads, in order to get as much information as possible about the behavior of the system in a limited period of time, usually when the building is unoccupied.

Long-term *performance monitoring* involves passive observation of the behavior of the systems of interest, waiting for operating conditions to change so as to reveal particular faults. Long-term monitoring requires the use of an Energy Management and Control System (EMCS) or a dedicated monitoring system. An approach that falls between these two extremes is to install special monitoring hardware, such as data loggers with internal memory, for a limited period, e.g. two weeks. This produces a ‘snapshot’ of building operation that can be analyzed to check for those faults that manifest themselves under the operating conditions encountered during the monitoring period.

The baseline that is used as a reference in performance monitoring can either be derived from design information and manufacturers' data or from performance data measured when the system is correctly operating. In the case of performance verification at commissioning time, only design information and manufacturers' data can be used, since the system has not yet operated.

Whole Building and Component-Level Diagnostics

Most current R&D in diagnostics can be categorized as either:

- **whole building level** – a 'top down' approach in which the main focus is on the energy consumption of the whole building and the major systems (chilled water, lighting, fans ...);
- **component level** – a 'bottom up' approach focussing on the operation of individual items of equipment and performance at the local loop level.

The ultimate aim is to link these two approaches.

At the whole building level, the baseline against which performance is compared can either be a database of the energy consumption of different buildings, such as CBECS (EIA 1995), or a simulation model, such as DOE-2, that defines the performance expected by the designer. Simulation models are starting to be used as the basis of performance contracts (DOE 1997). A third possibility is to fit an empirical model, such as an artificial neural network, to data collected when the building is deemed to be operating correctly.

At the component level, the baseline can be provided by design calculations, the manufacturer's performance data (corrected for allowed tolerances!) or performance measured *in situ*, either during commissioning or during a period of 'correct' operation.

Manual and Automated Procedures

Manual procedures rely on a human operator to initiate the collection and display of data and to analyze it, e.g. setting up an EMCS trend log and using it to trouble-shoot an occupant complaint.

Automated procedures use pre-programmed procedures to collect and analyze data on a routine basis and report conclusions without being prompted, e.g. monitoring mixing box temperatures and damper control signals at different operating points and diagnosing a leaking recirculation damper.

Semi-automated procedures use a combination of manual and automated methods, e.g. monitoring chiller efficiency and displaying a graph of kW/ton vs. tons if the kW/ton significantly exceeds the value expected for the current operating conditions.

Current commissioning procedures are almost all completely manual, in that they rely on a commissioning engineer or operator to perform the tests and analyze the results. There has been some R&D on automating performance verification procedures for use in commissioning (Haves, *et al.* 1996) but few commercial products appear to have been produced as yet. An automated commissioning tool is more complicated than a performance monitoring tool because it needs to be able to override the control system, either by forcing actuators or changing set-points, in order to exercise the equipment being tested. It also needs to include test sequences for each type of equipment for which it will be used.

Work on manual performance monitoring procedures for use during ‘normal’ operation has focused on graphical methods of displaying the measurements to the operator to allow efficient fault detection and then fault diagnosis. Displaying information to an operator is also an issue for automated tools, since such tools are unlikely to be accepted unless they can display the evidence and provide a justification for a fault detection and diagnosis. The main benefit of automated performance monitoring tools is that they can ‘pre-filter’ data from many points, avoiding the need for manual inspection of all the measurements from every point. By acting as ‘smart alarms’, automated performance monitoring tools have the potential to allow a building operator to spend less time keeping on top of performance and to allow remote operators and service companies to monitor multiple buildings efficiently. Ultimately, automated tools may be able to make reliable diagnoses and automatically contact service contractors and direct them to replace particular components.

EMCS or Dedicated Data Acquisition System?

Some diagnostic systems are intended to use data from the EMCS and some use dedicated data acquisition systems, which may either be installed permanently or temporarily. Issues include:

- purpose of monitoring (control performance, energy performance ...);
- number and type of sensors;
- sensor quality (accuracy, reliability, drift) ;
- sampling frequency and timestamp accuracy (15 minute intervals may be too long);
- data archiving (many EMCS’s have limited storage and retrieval capabilities);
- data visualization (many EMCS’s have limited manipulation and display capabilities);
- cost (installing a separate monitoring system adds significantly to first cost and also adds to maintenance costs).

Monitoring energy performance requires additional sensors that are not required for control purposes, e.g. electric power meters and fluid flow meters, and requires some sensors to be more accurate, e.g. chilled water flow and return temperature sensors when used to determine evaporator duty.

Review of Automated Fault Detection and Diagnosis Methods

Fault Detection

Fault *detection* is easier than fault *diagnosis* in that, for fault detection, it is only necessary to determine whether the performance is incorrect or unsatisfactory; knowledge of the different ways in which particular faults affect performance is not required. Two methods of representing the baseline for performance are:

- knowledge bases;
- quantitative models.

One common form of knowledge base is a set of rules produced by experts, e.g.

*IF the control valve is closed AND the supply fan is running AND the temperature rise across the heating coil is greater than the combined uncertainty of the sensor readings
THEN the operation is faulty.*

A knowledge base can be considered to be a qualitative model.

Quantitative models can take a number of forms, including:

- first principles ('textbook') models, e.g. effectiveness-NTU for coils;
- polynomial curve fits, e.g. fan curves or chiller performance maps;
- artificial neural networks.

Faults are detected with a quantitative model by using measured values of the control signals and some of the sensor signals as inputs to the model. The remaining sensor signals are then compared with the predictions of the model, as shown in Figure 1. Significant differences ('innovations') indicate the presence of a fault somewhere in the part of the system treated by the model.

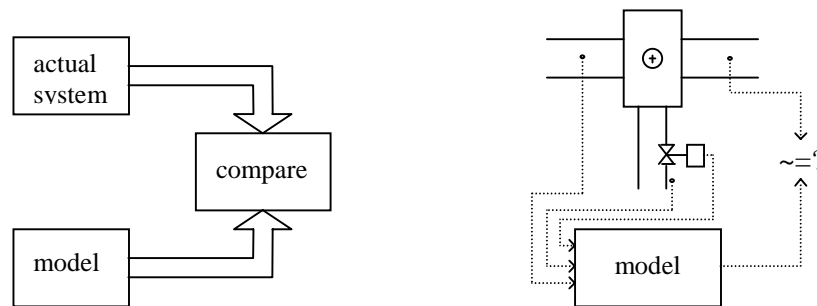


Figure 1: A model-based fault detection scheme. The left hand diagram shows the general principle and the right hand diagram shows the application to a heating coil, in which the inputs to the model are the inlet temperatures and the control signal and the output of the model is the outlet air temperature.

Empirical or ‘black box’ models, such as polynomial curve fits or neural networks are only as good as the data used to generate them. In particular, their accuracy usually falls off rapidly in regions of the operating space for which there are no training data. The main advantage of black box models is that they can be chosen to be linear in the parameters, making the process of estimating the parameters of the model both less demanding computationally and more robust. One advantage of physical models is that the prior knowledge that they embody improves their ability to extrapolate to regions of the operating space for which no training data are available. They also require fewer parameters for a given degree of model accuracy. A further feature of physical models is that the parameters correspond to physically meaningful quantities, which has two advantages:

1. Estimates of the values of the parameters can be obtained from design information and manufacturers’ data;
2. Abnormal values of particular parameters can be associated with the presence of particular faults.

Fault Diagnosis

Fault *diagnosis* is more difficult than fault *detection* because it requires knowledge of how the system behaves when faults are present. There are three main approaches:

1. Analysis of how the innovations (the differences between the predicted and the observed performance) vary with operating point;
2. Comparison of the actual behavior with the predictions of different models, each of which embodies a particular type of fault;
3. Estimation of the parameters of an on-line model that has been extended to treat particular faults.

One method of implementing the first approach is to use a rule-based classifier (Benouarets *et al.* 1994). The rules, which may be Boolean or fuzzy, can either be obtained from experts and then checked for consistency, completeness and correct implementation by testing using simulation. Alternatively, the rules can be generated using simulation. If black box models are used in the second approach, simulation may be the only way to generate training data, since it is not usually possible to obtain training data from real, faulty, systems. One exception is the case of factory-built standard products, such as chillers or VAV boxes, where the production of large numbers of identical units may justify the cost of generating training data for common faults. The third approach involves extending the model to represent both faulty and correct operation using physical principles and engineering knowledge; the estimated values for the parameters relating to the faulty behavior are then used for fault diagnosis (Buswell, *et al.* 1997).

Dynamic Behavior and Control Performance

In principle, reference models used in FDD should treat dynamic behavior as well as steady state behavior. However, the variations in operating point encountered in HVAC systems are often slow enough that most items of equipment can be regarded as being in a quasi-steady state, such that the error produced by using a static reference model is acceptably small. Static reference models are simpler to develop and configure; the dynamic behavior of HVAC equipment is often non-linear and poorly understood. Static models can be used for FDD if it is possible to determine when their predictions are valid, and when measurements can safely be used to estimate their parameters. Several *steady state detectors* have been produced by Annex 25 participants (Hyvarinen 1997).

Even though the dynamic response of most HVAC equipment is fast compared to typical rates of change of operating point, there are still circumstances in which the dynamic aspect of a system's controlled performance is important. Unstable operation, due to excessive controller gain, has a number of undesirable effects. These include excessive component wear, leading to premature failure and increased maintenance costs, and may also include discomfort or increased energy consumption. An overly damped response, due to low controller gain, may also lead to unacceptably sluggish performance in some situations.

An Integrated Approach to Performance Verification and Performance Monitoring

An approach to diagnostics that links design, commissioning and operation is to use a baseline derived from a combination of design data and manufacturers' performance data for performance verification. The data obtained during the performance verification tests is then used to 'fine-tune' this baseline for use in performance monitoring, as in the following procedure:

Performance Verification:

1. Configure the reference models of correct operation using design information and manufacturers' performance data.
2. Perform a predefined series of tests to verify correct performance.
3. If any tests fail, perform specific additional tests to determine the cause of failure.
4. Remedy defects diagnosed by open loop tests (e.g. replace mechanical equipment or retune controllers).
5. Repeat closed loop tests to verify system now performs correctly.

Performance Monitoring:

1. Use data collected during successful performance verification tests to refine the values of the parameters of the models of correct operation.
2. Use routine operating data to detect and, if possible, diagnose faults.
3. If a fault is detected but not adequately diagnosed, perform a performance verification test, typically during the next unoccupied period, to improve the diagnosis of the fault.

4. Remedy the fault.
5. Confirm that the fault has been correctly remedied by repeating the performance verification test.

Evaluation of Diagnostic Systems

Diagnostic systems can be evaluated by assessing the extent to which they supply reliable, needed, information at low cost. In particular, a good diagnostic system should:

1. Have sufficient sensitivity to detect faults before they become problems (e.g. cause complaints or damage);
2. Provide a useful diagnosis by:
 - a) Localizing the fault (which building, which floor, which AHU, which coil, is it in the valve or the actuator...?);
 - b) Recommending action (which trade to call: mechanical, electrical, controls?), is action required immediately (what is the expected cost of delay?);
3. Have a low false alarm rate;
4. Have a low first cost and low maintenance costs;
5. Be easy to learn and use.

Acknowledgement

This work was supported by the California Institute for Energy Efficiency (CIEE) and the Assistant Secretary for Energy Efficiency and Renewable Energy, Office of Building Technology, State and Community Programs of the U.S. Department of Energy under Contract No. DE-AC03-76SF00098. CIEE is a research unit of the University of California.

References

- Benouarets, M., A.L. Dexter, R.S. Fargus, P. Haves, T.I. Salsbury and J.A. Wright. 1994. "Model-Based Approaches to Fault Detection and Diagnosis in Air-Conditioning Systems," *Proceedings of System Simulation in Buildings '94*, Liege, Belgium.
- Buswell, R.A., P. Haves and T.I. Salsbury. 1997. "A Model-Based Approach to the Commissioning of HVAC Systems," *Proceedings of CLIMA 2000*, Brussels.
- DOE 1997, *International Performance Measurement and Verification Protocol*, US Department of Energy, DOE/EE-0157
- EIA 1995. *Commercial Buildings Energy Consumption and Expenditures 1992*, Commercial Buildings Energy Consumption Survey, Energy Information Administration. U.S. Department of Energy
- Haves, P., D.R. Jorgensen, T.I. Salsbury and A.L. Dexter. 1996. "Development and Testing of a Prototype Tool for HVAC Control System Commissioning," *ASHRAE Transactions*, Vol. 102, Pt 1.
- Hyvarinen, J. 1997. IEA Annex 25 Final Report. VTT, Espoo, Finland.

